WE CLAIM:

1. A method of transmitting information through a data switching apparatus connected to a plurality of input line end devices (ILE_{00}, ILE_{01}, ILE_{0m}) and output line end devices (ELE_{00}, ELE_{01}, ELE_{0m}), said input line end devices (ILE_{00}, ILE_{01}, ILE_{0m}) transmitting protocol information packets to the data switch for transmission to specific output line end devices (ELE_{00}, ELE_{01}, ELE_{0m}),

the data switching apparatus comprising a plurality of input traffic manager units (ITM_0, ITM_1, ITM_n) a plurality of output traffic manager units (ETM_0, ETM_1, ... ETM_n) and a data switch (SW), the data switch (SW) comprising a plurality of input routers (SRI_0, SRI_1, ... SRI_p), a plurality of output routers (SRE_0, SRE_1, ... SRE_p), and a memory-less cyclic switch fabric (SCM), and a switch controller (SM), said switch fabric being controlled by said switch controller (SM), said input traffic manager units (ITM_0, ITM_1, ITM_n) being connected to one or more of said input line end devices (ILE_{00}, ILE_{01}, ILE_{0m}), and said output traffic manager units (ETM_0, ETM_1, ... ETM_n) being connected to one or more of said output line end devices (ELE_{00}, ELE_{01}, ELE_{0m}),

each input traffic manager unit (ITM_0, ITM_1, ITM_n) being arranged to convert the protocol information packets it receives from the respective input line end devices (ILE_{00}, ILE_{01}, ILE_{0m}) into fixed length cells having a header (UH), said header (UH) indicating the output traffic manager unit (ETM_{01} ETM_1, ... ETM_n) connected to the output line end device (ELE_{00}, ELE_{01}, ELE_{0m}) to which the cell should be sent,

each input router (SRI_{01}, SRI_1, ... SRI_p) being arranged to receive cells from a respective group of said input traffic manager units (ITM_0, ITM_1, ... ITM_n), and to

maintain virtual output queues for each output traffic manager unit (ETM_0, ETM_1, ... ETM_n);

each output router (SRE_0, SRE_1, ...SRE_p) being arranged to transmit cells to a respective group of said output traffic manager units (ETM_0, ETM_1, ... ETM_n);

the method comprising, on the arrival of a cell from an input traffic manager unit (ITM_0, ITM_1, ITM_n) the input router (SRI_0, SRI_1, ... SRI_p) examining the cell header (UH), placing it in a virtual output queue for the output traffic manager unit (ETM_0, ETM_1, ...ETM_n) indicated by the cell header (UH), generating a transfer request (RFT) including the address of the output traffic manager unit (ETM_0, ETM_1, ... ETM_n) indicated by the header (UH) of that cell, and passing said request (RFT) to the switch controller (SM),

characterized in that:

said cell headers (UH) include message priority information, and said transfer requests (RFT) include a priority code;

the switch fabric (SCM) is controlled by the switch controller (SKI) to connect ones of said input routers (SRI_0, SRI_1, ... SRI_p) to ones of said output routers (SRE_0, SRE_1, ... SRE_p);

the switch controller (SM) schedules the passage of the cells across the switch fabric (SCM) at each switch cycle, by using a first arbitration process to select which of said input routers (SRI_0, SRI_1, ... SRI_p) to connect to which of said output routers (SRE_0, SRE_1, ... SRE_p), and controls the switch fabric to connect the selected input routers (SRI_0, SRI_1, ... SRI_p) to the corresponding selected output routers (SRE_0,

SRE_1, ... SRE_p); and

upon it being determined that a given input router (SRI_0, SRI_1, ... SRI_p) is to be connected to a given output router (SRE_0, SRE_1, ... SRE_p):

that given input router (SRI_0, SRI_1, ... SRI_p) performs a second arbitration process to select a single virtual output queue, from among the virtual output queues for the output traffic manager units (ETM_0, ETM_1, ... ETM_n) to which the given output router (SRE_0, SRE_1, ... SRE_p) sends cells, and transmits the cell at the head of the selected virtual output queue across the switch fabric (SCM) to the given output router (SRE_0, SRE_1, ... SRE_p),

and the given output router (SRE_0, SRE_1, ... SRE_p) transmits the cell to the output traffic manager unit (ETM_0, ETM_1, ... ETM_n) indicated by the cell header (UH).

2. A method according to claim 1 in which each input router (SRI_0, SRI_1, ... SRI_p) maintains a virtual output queue for each output traffic manager unit (ETM_0, ETM_1,... ETM_n) and priority level, and upon receipt of a cell the input router (SRI_0, SRI_1, ... SRI_p) places the cell in the virtual output queue for the priority and output traffic manager unit (ETM_0, ETM_1, ... ETM_n) indicated by the cell header (UH).

3. A method according to claim 1 or 2 in which each output router (SRE_0, SRE_1, ... SRE_p) maintains an output queue for each of the group of output manager units (ETM_0, ETM_1, ... ETM_n) to which it transmits cells.

4. A method according to any preceding claim in which each input router (SRI_0, SRI_1, ... SRI_p) maintains an input buffer for each of the group of input traffic manager units (ITM_0, ITM_1, ... ITM_n) from which it receives signals.

5. A method according to any preceding claim in which said second arbitration process performed by the given input router (SRI_0, SRI_1, ... SRI_p) is a weighted round-robin arbitration process based upon: the length of said output virtual queues of the given input router (SRI_0, SRI_1, ... SRI_p); an aggregate queue packet urgency; and a backpressure from said output traffic manager units (ETM_0, ETM_1, ... ETM_n).

6. A method according to any preceding claim in which the first arbitration process selects which input routers (SRI_0, SRI_1, ... SRI_p) and output routers (SRE_0, SRE_1, ... SRE_p) to connect, to maximise the number of said requests (RIFT) which can be satisfied.

7. A data switching apparatus for connection to a plurality of input line end devices (ILE_{00}, ILE_{01}, ILE_{0m}) and output line end devices (ELE_{00}, ELE_{01}, ELE_{0m}) to transmit protocol information packets received from said input line end devices (ILE_{00}, ILE_{01}, ILE_{0m}) to specific output line end devices (ELE_{00}, ELE_{01}, ELE_{0m}),

the data switching apparatus comprising a plurality of input traffic manager units (ITM_0, ITM_1, ITM_n), a plurality of output traffic manager units (ETM_0, ETM_1, ... ETM_n) and a data switch (SW), the data switch (SW) comprising a plurality of input routers (SRI_0, SRI_1, ... SRI_p), a plurality of output routers (SRE_0, SRE_1, ... SRE_p), a

memory-less cyclic switch fabric (SF), and a switch controller (SM), said switch fabric being controlled by said switch controller, each of said input traffic manager units (ITM_0, ITM_1, ITM_n) being for connection to one or more of said input line end devices (ILE_{00}, ILE_{01}, ILE_{0m}), and each of said output traffic manager units (ETM_0, ETM_1,... ETM_n) being for connection to one or more of said output line end devices (ELE_{00}, ELE_{01}, ELE_{0m}),

each input traffic manager unit (ITM_0, ITM_1, ITM_n) being arranged to convert the protocol information packets it receives from the respective input line end devices (ILE_{00}, ILE_{01}, ILE_{0m}) into fixed length cells having a cell header (UH), said cell header (UH) indicating the output traffic manager unit (ETM_{00}, ETM_{01} ... ETM_n) connected to the output line end device (ELE_{00}, ELE_{01}, ELE_{0m}) to which the cell should be sent,

each of the input routers (SRI_0, SRI_1, ... SRI_p) being arranged to receive cells from a respective group of said input traffic manager units (ITM_0, ITM_1, ITM_n), to maintain a set of virtual output queues for each output traffic manager unit (ETM_0, ETM_1, ... ETM_n), and, on the arrival of a cell from an input traffic manager unit (ITM_0, ITM_1, ITM_n), to examine the cell header (UH), to place it in a virtual output queue for the output traffic manager unit (ETM_0, ETM_1, ... ETM_n) indicated by the cell header (UH), to generate a transfer request (RFT) including the address of the output traffic manager unit (ETM_0, ETM_1,...ETM_n) indicated by the header (UH) of that cell, and to pass said request (RIFT) to the switch controller,

each output router (SRE_0, SRE_1, ... SRE_p) being connected to a

respective group of said output traffic manager units (ETM_0, ETM_1, ... ETM_n);

characterized in that:

each output router (SRE_0, SRE_1, ... SRE_p) is arranged, upon receipt of a cell having a header (UH) which indicates one of that group of output traffic manager units (ETM_0, ETM_1, ... ETM_n), to transmit the cell to that indicated output traffic manager unit (ETM_0, ETM_1, ... ETM_n);

said input traffic manager units (ITM_0, ITM_1, ITM_n) are arranged to include message priority information in said cell headers (UH), and said input routers (SRI_0, SRI_1, ... SRI_p) are arranged to include a priority code in said transfer requests (RFT);

the switch fabric (SCM) is arranged, under the control of the switch controller (SM), to connect ones of said input routers (SRI_0, SRI_1, ... SRI_p) to ones of said output routers (SRE_0, SRE_1, ... SRE_p);

the switch controller (SM) is arranged to schedule the passage of the cells across the switch fabric at each switch cycle, by using a first arbitration process to select which of said input routers (SRI_0, SRI_1, ... SRI_p) to connect to which of said output routers (SRE_0, SRE_1, ... SRE_p), and control the switch fabric to connect the selected input routers (SRI_0, SRI_1, ... SRI_p) to the corresponding selected output routers (SRE_0, SRE_1, ... SRE_p); and

each input router (SRI_0, SRI_1, ... SRI_p) is arranged, upon it being determined that that input router (SRI_0, SRI_1, ... SRI_p) is to be connected to a given output router (SRE_0, SRE_1, ... SRE_p), to perform a second arbitration process to select a single

virtual output queue from among the virtual output queues for the output traffic manager units (ETM_0, ETM_1, ... ETM_n) to which the given output router (SRE_0, SRE_1, ... SRE_p) is connected, and to transmit the cell at the head of the selected virtual output queue across the switch fabric (SF) to the given output router (SRE_0, SRE_1, ... SRE_p).

8. A data switching apparatus according to claim 7 in which, each input router (SRI_0, SRI_1, ... SRI_p), is arranged to maintain a virtual output queue for each output traffic manager unit (ETM_0, ETM_1, ... ETM_n) and priority level, and the input router (SRI_0 SRI_1 ... SRI_p)is arranged to place a received cell in the virtual output queue for the priority and output traffic manager unit (ETM_0, ETM_1, ... ETM_n) indicated by the cell header (UH).

9. A data switching apparatus according to claim 7 or 8 in which each output router (SRE_0, SRE_1, ... SRE_p) is arranged to maintain an output queue for each of the group of output manager units (ETM_0, ETM_1, ... ETM_n) to which it can send cell.

10. A data switching apparatus according to any of claims 7 to 9 in which each input router (SRI_0, SRI_1, ... SRI_p) is arranged to maintain an input buffer for each of the group of input traffic manager units (ITM_0, ITM_1, ITM_n) from which it receives signals.

11. A data switching apparatus according to any of claims 7 to 10 in which said second arbitration process is a weighted round-robin arbitration process based upon: the length of said output virtual queues of the given input router (SRI_0, SRI_1, ...

SRI_p); an aggregate queue packet urgency; and a backpressure from said output traffic manager units (ETM_0, ETM_1, ... ETM_n).

12. A data switching apparatus according to any of claims 7 to 11 in which the first arbitration process selects which input routers (SRI_0, SRI_1, ... SRI_p) and output routers (SRE_0, SRE_1, ... SRE_p) to connect, to maximise the number of said requests (RIFT) which can be satisfied.